You requested this document:

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

1. **Modular synchronization in distributed, multiversion databases: version control control**
Agrawal, D.; Sengupta, S.;
Knowledge and Data Engineering, IEEE Transactions on
Volume 5, Issue 1, Feb. 1993 Page(s):126 - 137
Abstract:

A version control mechanism is proposed that enhances the modularity and extensibilit concurrency control algorithms. The multiversion algorithms are decoupled into two cor control and concurrency control. This permits modular development of multiversion pro simplifies the task of proving the correctness of these protocols. A set of procedures fo described that defines the interface with the version control component. It is shown tha interface can be used by the database actions of both two-phase locking and time-stan control protocols to access multiversion data. An interesting feature of the framework is execution of read-only transactions becomes completely independent of the underlying control implementation. Unlike other multiversion algorithms, read-only transactions in not modify any version-related information, and therefore do not interfere with the exec transactions. The extension of the multiversion algorithms to a distributed environment simple

Abstract | Full Text: PDF(1208 KB) **IEEE JNL**

Replication: DB2, Oracle, or Sybase?

Full text Pdf (727 KB)

Source **ACM SIGMOD Record** archive
Volume 24 , Issue 4 (December 1995) table of contents
Pages: 95 - 101
Year of Publication: 1995
ISSN:0163-5808

Author Doug Stacey Comdisco Inc., Rosemont, IL

Publisher ACM New York, NY, USA

↑ ABSTRACT

Is replication salvation or the devil in disguise? Here's what three implementations tell us

↑ CITED BY *4*

Stéphane Gançarski , Hubert Naacke , Esther Pacitti , Patrick Valduriez, The leganet system: Freshness-aware transaction routing in a database cluster, Information Systems, v.32 n.2, p.320-343, April, 2007

Parvathi Chundi , Daniel J. Rosenkrantz , S. S. Ravi, Multi-site distributed database transactions utilizing deferred update, Proceedings of the 1997 ACM symposium on Applied computing, p.118-122, April 1997, San Jose, California, United States

Fernando Pedone , Rachid Guerraoui , André Schiper, The Database State Machine Approach, Distributed and Parallel Databases, v.14 n.1, p.71-98, July 2003

Richard Hull, Managing semantic heterogeneity in databases: a theoretical prospective, Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems, p.51-61, May 11-15, 1997, Tucson, Arizona, United States

↑ INDEX TERMS

Primary Classification:
H. Information Systems
↳ **H.2** DATABASE MANAGEMENT

Additional Classification:
H. Information Systems
↳ **H.2** DATABASE MANAGEMENT
↳ **H.2.3** Languages
↳ **Nouns:** SYBASE; DB2
↳ **H.2.4** Systems
↳ **Nouns:** ORACLE; SYBASE

General Terms:
Design, Management, Performance, Theory

↑ **Peer to Peer - Readers of this Article have also read:**

- Data structures for quadtree approximation and compression **Communications of the ACM** 28, 9
  Hanan Samet
- A hierarchical single-key-lock access control using the Chinese remainder theorem **Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computing**
  Kim S. Lee , Huizhu Lu , D. D. Fisher
- The GemStone object database management system **Communications of the ACM** 34, 10
  Paul Butterworth , Allen Otis , Jacob Stein
- Putting innovation to work: adoption strategies for multimedia communication systems **Communications of the ACM** 34, 12
  Ellen Francik , Susan Ehrlich Rudman , Donna Cooper , Stephen Levine
- An intelligent component database for behavioral synthesis **Proceedings of the 27th ACM/IEEE conference on Design automation**
  Gwo-Dong Chen , Daniel D. Gajski

Google Scholar BETA

Web Images Video News Maps more »

synchronization databases log 1991 - 1995 Search Ad Sc Sc

Scholar **All articles - Recent articles** Results **1** - **10** of about **1,760** for **synchronization databases l**

All Results
C Mohan
R Golding
M Seltzer
H Garcia-Molin...
E Rahm

Concurrently applying redo records to backup database in a **log** sequence using single queue server ... - all 2 versions »
C Mohan, RL Obermarck, RK Treiber - US Patent 5,170,480, 1992 - Google Patents
... transaction consistency, active **log** data processing ... alternate procedures for maintaining
tracking **databases**. Foremost, **synchronization** is not required between ...
Cited by 57 - Related Articles - Web Search

[PDF] VINO: An Integrated Platform for Operating System and Database Research - all 5 versions »
C Small, M Seltzer - Cambridge, MA, Tech. Rep. TR-30-94, 1994 - dogfish.org
... the Berkeley Fast Filesystem, and the Berkeley **Log**-structured Filesys ... under- lying system support for a **database** managementsystem ... 5.1 **Synchronization** Primitives ...

Cited by 37 - Related Articles - View as HTML - Web Search

Main memory database systems: an overview - all 7 versions »
H Garcia-Molina, K Salem - Knowledge and Data Engineering, IEEE Transactions on, 1992 - ieeexplore.ieee.org
... Checkpointing brings the disk resident copy of the **database** more up-to-date, thereby eliminating the need for the least recent **log** entries. ...
Cited by 225 - Related Articles - Web Search - Library Search

[PDF] Transaction Support in a **Log**-Structured File System - all 6 versions »
M Seltzer - Proceedings of the Ninth International Conference on Data ..., 1993 - eecs.harvard.edu
... 1] show how to implement user **synchronization** quickly on ... sequential scan, the structure
of the **database** will be ... realize a benefit from the **log**- structured file ...
Cited by 18 - Related Articles - View as HTML - Web Search

Modular **synchronization** in distributed, multiversion **databases**: version control and concurrency ... - all 7 versions »
D Agrawal, S Sengupta - IEEE Transactions on Knowledge and Data Engineering, 1993 - doi.ieeecomputersociety.org
... IEEE **Log** Number 9205827. ... Page 2. AGRAWAL AND SENGUFTA: MODULAR **SYNCHRONIZATION** IN DISTRIBUTED MULTIVERSION **DATABASES** 127 ...
Cited by 17 - Related Articles - Web Search - Library Search

Empirical performance evaluation of concurrency and coherency control protocols for database sharing ... - all 5 versions »
E Rahm - ACM Transactions on Database Systems (TODS), 1993 - portal.acm.org
... Transaction routing I I '1 '2 . . . 'N &b D (**database**, **log** files) (a) (b) Fig. 1. DB-sharing versus DB-partitioning (a) DB-sharing (b) DB-partitioning, ...
Cited by 49 - Related Articles - Web Search

Multimedia **synchronization** protocols for broadband integratedservices -

all 5 versions »
TDC Little, A Ghafoor - Selected Areas in Communications, IEEE Journal on, 1991 - ieeexplore.ieee.org
... of audio, video, text, and numeric data originating from **databases** or live ... dependent data is the need to provide **synchronization** of data ... IEEE **Log** Number 9102860 ...
Cited by 163 - Related Articles - Web Search

Implementing agent coordination for workflow management systemsusing active database systems - all 2 versions »
C Bussler, S Jablonski - Research Issues in Data Engineering, 1994. Active Database ..., 1994 - ieeexplore.ieee.org
... As well as activity specifications and history (**log**) information, **synchronization** policies must be stored in **databases** in order to maintain consistency. ...
Cited by 56 - Related Articles - Web Search

Context-specific **synchronization** for atomic data types in object-based **databases** - all 5 versions »
MH Wong, D Agrawal - Theoretical Computer Science, 1995 - Elsevier
... Theoretical Computer Science Context-specific **synchronization** for atomic data types in object-based **databases** Man Hon Wong", Divyakant Agrawal b.* ...
Cited by 4 - Related Articles - Web Search

Lightweight recoverable virtual memory - all 52 versions »
M Satyanarayanan, HH Mashburn, P Kumar, DC Steere, ... - ACM Transactions on Computer Systems (TOCS), 1994 - portal.acm.org
... range of applications from distributed file systems and **databases** to object- ... choice of logging, **synchronization**, and transaction commitment strategies. ...
Cited by 130 - Related Articles - Web Search - Library Search

Goooooooooogle ▶
Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

synchronization databases log Search

Google Home - About Google - About Google Scholar

Management of a remote backup copy for disaster recovery

Full text Pdf (2.48 MB)

Source **ACM Transactions on Database Systems (TODS)** archive
Volume 16 , Issue 2 (June 1991) table of contents
Pages: 338 - 368
Year of Publication: 1991
ISSN:0362-5915

Authors Richard P. King IBM T. J. Watson Research Center, Yorktown Heights, NY
Nagui Halim IBM T. J. Watson Research Center, Yorktown Heights, NY
Hector Garcia-Molina Princeton Univ., Princeton, NJ
Christos A. Polyzois P rinceton Univ., Princeton, NJ

Publisher ACM New York, NY, USA

↑ ABSTRACT

A remote backup database system tracks the state of a primary system, taking over transaction processing when disaster hits the primary site. The primary and backup sites are physically isolated so that failures at one site are unlikely to propogate to the other. For correctness, the execution schedule at the backup must be equivalent to that at the primary. When the primary and backup sites contain a single processor, it is easy to achieve this property. However, this is harder to do when each site contains multiple processors and sites are connected via multiple communication lines. We present an efficient transaction processing mechanism for multiprocessor systems that guarantees this and other important properties. We also present a database initialization algorithm that copies the database to a backup site while transactions are being processed.

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

1 ~kGRAWAL, R. A parallel logging algorithm for multiprocessor database machines. In Proceedings of the 4th Internahonal Workshop on Database Machines. Springer, New York, 1985.

2 Philip A. Bernstein , Vassco Hadzilacos , Nathan Goodman, Concurrency control and

recovery in database systems, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, 1987

3 BURKES, D., AND TREIBER, K. Design approaches for real time recovery. Presentation at the Third International Workshop on High Performance Transaction Systems (Pacific Grove, Calif., Sept. 1989).

4 CRUS, R.A. Data recovery in IBM Database 2. IBM Syst. J. 23, 2 (1984), 178-188.

5 FINKELSTEIN, W., AND CAPPI, M. Experiences with large networks of computers. Presenta- Lion at the International Workshop on High Performance Transaction Systems (Pacific Grove, Calif., Sept. 1985).

6 GARCIA-MOLINA, H., AND ABBOTT, R. K. Reliable distributed database management. In Proceedings of the IEEE, Special Issue on D~stributed Database Systems (May 1987), 601-620.

7 GRAY, J. N., AND ANDERTON, M. Distributed computer systems: Four case studies. In Proceedings of the IEEE, Special Issue on Distributed Database Systems (May 1987), 719-726.

8 GRAY, J.N. Why do computers stop and what can be done about it? Presentation at the Fifth Symposium on Reliability in Distributed Software and Database Systems (Los Angeles, Calif., Jan. 1986).

9 Jim Gray, Notes on Data Base Operating Systems, Operating Systems, An Advanced Course, p.393-481, January 1978

10 GRAY, J. N., AND t~EUTER, A. Transaction processing. Course Notes from CS~,445 StanforJ Spring Term, 1988.

11 Henry F. Korth , Abraham Silberschatz, Database system concepts, McGraw-Hill, Inc., New York, NY, 1986

12 IBM, IMS/VS Extended Recovery Facility (XRF)' General Informatzon. Doc. GG24-3150, March 1987

13 LYON, J Design considerations in replicated database systems for dmaster protectmn. IEEE Compcon, 1988

14 Patrick E. O'Neil, The Escrow transactional method, ACM Transactions on Database Systems (TODS), v.11 n.4, p.405-430, Dec. 1986 [doi>10.1145/7239.7265]

15 Daniel J. Rosenkrantz, Dynamic database dumping, Proceedings of the 1978 ACM SIGMOD international conference on management of data, May 31-June 02, 1978, Austin, Texas [doi>10.1145/509252.509257]

16 Dale Skeen, Nonblocking commit protocols, Proceedings of the 1981 ACM SIGMOD international conference on Management of data, April 29-May 01, 1981, Ann Arbor, Michigan [doi>10.1145/582318.582339]

17 Richard D. Schlichting , Fred B. Schneider, Fail-stop processors: an approach to designing fault-tolerant computing systems, ACM Transactions on Computer Systems (TOCS), v.1 n.3, p.222-238, August 1983 [doi>10.1145/357369.357371]

18 Tandem Computers Remote Duplicate Database Factllty (RDF) System Management Manual. March 1987

19 A. S. Tanenbaum, Computer networks, Prentice-Hall, Inc., Upper Saddle River, NJ, 1988

↑ CITED BY 9

- Sharad Mehrotra , Kexiang Hu , Simon Kaplan, Dealing with partial failures in multiple processor primary-backup systems, Proceedings of the sixth international conference on Information and knowledge management, p.371-378, November 10-14, 1997, Las Vegas, Nevada, United States
- Christos A. Polyzois , Hector Garcia-Molina, Evaluation of remote backup algorithms for transaction processing systems, ACM SIGMOD Record, v.21 n.2, p.246-255, June 1, 1992
- Manhoi Choy , Hong Va Leong , Man Hon Wong, Disaster recovery techniques for database systems, Communications of the ACM, v.43 n.11es, Nov. 2000
- David B. Lomet, High speed on-line backup when using logical log operations, ACM SIGMOD Record, v.29 n.2, p.34-45, June 2000
- Christos A. Polyzois , Héctor García-Molina, Evaluation of remote backup algorithms for transaction-processing systems, ACM Transactions on Database Systems (TODS), v.19 n.3, p.423-449, Sept. 1994

  Subhash Bhalla , Stuart E. Madnick, Parallel on-the-fly reading of an entire database copy, Practical parallel computing, Nova Science Publishers, Inc., Commack, NY, 2001

  Subhash Bhalla , Stuart E. Madnick, Asynchronous Backup and Initialization of a Database Server for Replicated Database Systems, The Journal of Supercomputing, v.27 n.1, p.69-89, January 2004

  Gerhard Weikum , Christof Hasse, Multi-level transaction management for complex objects: implementation, performance, parallelism, The VLDB Journal — The International Journal on Very Large Data Bases, v.2 n.4, p.407-454, October 1993
- Erhard Rahm, Empirical performance evaluation of concurrency and coherency control protocols for database sharing systems, ACM Transactions on Database Systems (TODS), v.18 n.2, p.333-377, June 1993

↑ INDEX TERMS

Primary Classification:

C. Computer Systems Organization
- **C.2** COMPUTER-COMMUNICATION NETWORKS
  - **C.2.4** Distributed Systems
    - **Subjects:** Distributed databases

Additional Classification:

D. Software
- **D.4** OPERATING SYSTEMS
  - **D.4.5** Reliability
    - **Subjects:** Backup procedures

H. Information Systems
- **H.2** DATABASE MANAGEMENT
  - **H.2.4** Systems
    - **Subjects:** Transaction processing
  - **H.2.7** Database Administration

Subjects: Logging and recovery

General Terms:
Algorithms, Reliability

Keywords:
database initialization, hot spare, hot standby, remote backup

REVIEW

"Peter John Trueman"

Two multicomputer systems are linked by multiple communication paths, and each has a partitioned relational database. The primary system processes the database transactions, and the other acts as a hot-standby that will take over the processing in the event of the primary system failing. This paper describes the decentralized algorithm used to ensure that the backup system's database is up to date and consistent. The algorithm is intended for use in applications that have a lot of transactions and need a quick response time. So, rather than use an expensive two-phase commit protocol to ensure that a transaction atomically updates both systems, the transaction simply commits and then propagates to the backup. This means that when the primary fails, some committed transactions may be lost and others may have to be discarded to maintain consistency; for example, if the transaction that created a bank account is lost, updating that account would be wrong. This risk is deemed to be economically acceptable given the performance requirement; if the consequences of a lost transaction are high, however, it is possible to use an atomic commit over both systems. After some good background and introductory sections, the paper defines what it means for the backup system to be consistent, and then describes how the backup system is initialized and updated, and argues that this process results in a consistent backup. Although the paper is lengthy, it is not too long; it is well structured and clear. The reader who wants just as much information as can be easily remembered can read the first ten pages, leaving the rest for the serious student. The paper is well worth reading. *Online Computing Reviews Service*

Collaborative Colleagues:

Hector Garcia-Molina:	Soraya Abad	Neil Daswani	Steven K. Ketchpel	Gerhard Rossbach
	Robert Abbott	Susan Davidson	Steven P. Ketchpel	Nick Roussopoulos
	Robert K. Abbott	Susan B. Davidson	Steven Paul Ketchpel	Kenneth Salem
	Serge Abiteboul	Umesh Dayal	Johannes Klein	Hans-Jörg Schek
	Brad Adelberg	David De Witt	Karl Kleissner	Mario Schlosser
	Rafael Alonso	David DeWitt	Henry F. Korth	Mario T. Schlosser
	Arvind Arasu	Hrishikesh Deshpande	Wilburt Labio	Amit P. Sheth
	Michelle Baldonado	Min Fang	Wilburt J Labio	Narayanan Shivakumar
	Daniel Barbará	Prasanna Ganesan	Wilburt J. Labio	Dale Skeen
	Daniel Barbará-Millá	Roy Goldman	Wilburt Juan Labio	Yee Jiun Song
	Pavel Berkhin	David Gordon	Carl Lagoze	AnneMarie Spauster
	Phil Bernstein	Adrian Graham	Wang Lam	Annemarie Spauster
	David Bloom	Luis Gravano	Michael Lesk	Carl Staelin
	James Brady	Zoltán Gyöngyi	Elizabeth Liddy	Mike Stonebraker
	Onn Brandman	Zoltan Gyongi	Bruce Lindsay	Qi Sun
	Yuri Breitbart	Zoltan Gyongyi	David B. Lomet	Qixiang Sun
	Orkut Buyukkokten	Robert B. Stefano Ceri	Dave Maier	Jens Tellefsen
			Sergio Marti	Stanford University
			Sergey Melink	Murty Valiveti
	Chen Chang		Sergey Melnik	Suresh Venkatasubramanian

	Chen-Chuan Chang Edward Yjhuei Chang Kevin Chang Kevin C Chang Sudarshan S. Chawathe Sudarshan Sudhir Chawathe Chandra Chekuri Junghoo Cho Jae Chung Tyson Condie Brian F. Cooper Brian Frank Cooper Steve Cousins Steve B. Cousins Arturo Crespo Yingwei Cui	Hagmann Nagui Halim Susumu Harada Scott W. Hassan Taher Haveliwala Jerry Held Joe Hellerstein Diane Hillmann Jun Hirai Mei Hsu Matthew Jacobsen Stephen Johnson Oliver Kaljuvee Sepandar Kamvar Sepandar D. Kamvar Ben Kao Benjamin Kao Seth Katz Jack Kent	Subhasish Mitra Bob Mungamuru Lawrence Page Yannis Papakonstantinou Jan Pedersen Hamid Pirahesh Frank Pittelli Christos A. Polyzios Christos A. Polyzois Daryl Porter Dallan Quass Martin Röscheisen Sriram Raghavan Gerard Rodriguez-Mula Martin Roscheisen	Patrick Vinograd Rosemary Walsh Michelle Q. Wang Baldonado QianYing Wang Stuart Weibel Gary Wesley Tak W Yan Tak W. Yan Tak Woon Yan Beverly Yang Ron B. Yeh Ramana Yerneni Ramana Venkata Yerneni Yue Zhuge Jaco Zijlstra
Nagui Halim:	Kay S. Anderson Eric Bouillet Hector Garcia-Molina Sheng Ma Christos A. Polyzois David J. Taylor			
Christos A. Polyzois:	Anupam Bhide Daniel M. Dias Héctor García-Molina Hector Garcia-Molina Robert B. Hagmann Nagui Halim François Ménard K. Hal Purdy Henning Schulzrinne David Shrader	Henry Sinnreich Ping-Fai Yang		

↑ **Peer to Peer - Readers of this Article have also read:**

- Inferring constraints from multiple snapshots **ACM Transactions on Graphics (TOG)** 12, 4
  David Kurlander , Steven Feiner

- Data structures for quadtree approximation and compression **Communications of the ACM** 28, 9
  Hanan Samet

- A hierarchical single-key-lock access control using the Chinese remainder theorem **Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computing**
  Kim S. Lee , Huizhu Lu , D. D. Fisher

- The GemStone object database management system **Communications of the ACM** 34, 10
  Paul Butterworth , Allen Otis , Jacob Stein

- An intelligent component database for behavioral synthesis **Proceedings of the 27th ACM/IEEE conference on Design automation**
  Gwo-Dong Chen , Daniel D. Gajski

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L29	166	database with synchroniz$4 with log and database	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/07 15:34
L30	28	database with synchroniz$4 with log same (single or central or shared) and database	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/07 15:36